

Market-based Grid Computing - Survey Results

This paper summarizes the white papers and the current implementation being used on some of the Grid sites on the Distributed Accounting system. It's evident that the traditional accounting method, where users are required to have an account on the system to use the resources, needs new mechanism for the Grid computing to scale. There are few white papers posted on the GRID Forum Accounts working group(<http://www.nas.nasa.gov/~thigpen/accounts-wg/>) which addresses these issues (problem of maintaining thousands of user accounts on the system, etc.) and introduces the new method of allowing authorized users to use the resources on the local systems without requiring users to have the local account on the system.

There are already two centers, Polish National Super Computer and University of Manchester of England, which uses account-less model to allow users to run the job using the template accounts to bind to the real user during the lifetime of the job execution. Following white papers listed under GF Accounts-wg and the URL(<http://www.nas.nasa.gov/~thigpen/accounts-wg/>) has detail informations on the account-less model using the template accounts.

- Account Allocation on the Grid
- Simplifying Administration and Management Processes in the Polish National Clusters
- <http://www.hep.grid.ac.uk/gridmapdir/>

Grid Computing Tools and Accounts model

The following Table 1 summarizes some of the middleware software that are currently been deployed on the Distributed Computing systems with the reference to the User Accounting Model currently being used with the middleware software.

Table 1:

Middleware	User Account Model	Authentication Method	Authorization Method	Advantage/disadvantage
Condor	Uses one UID ("nobody") to execute jobs for users that don't have an local account in a Condor flock.	Standard UNIX user authentication	Existing account on Condor front-end systems	Advantage: <ul style="list-style-type: none">• Easy to maintain user accounts Disadvantage: <ul style="list-style-type: none">• Difficult to distinguish users• interference in scratch space• lack of effective account-ability• problem in message queues and shared memory for the process in UNIX.

Table 1:

Middleware	User Account Model	Authentication Method	Authorization Method	Advantage/disadvantage
PUNCH	Users are represented by logical user accounts within a single physical account with the ability to use a set of dynamic account bindings for system calls.	Standard UNIX user authentication	N/A	Advantage: <ul style="list-style-type: none"> • Easy to maintain user accounts Disadvantage: <ul style="list-style-type: none"> • Difficult to distinguish users • interference in scratch space • lack of effective accountability • problem in message queues and shared memory for the process in UNIX.
Globus	Grid user has individual local accounts with the entry in the Grid mapfile.	X.509 Certificate Grid mapfile	Entry in the gridmap file with the corresponding local unix account.	Advantage: <ul style="list-style-type: none"> • Easy to identify each Grid user for accounting, accountability Disadvantage: <ul style="list-style-type: none"> • Problem of maintaining user accounts and Grid mapfile.

The Table 2 summarizes the current implementation of the account-less model on some sites and the white papers on the Distributed Computing model.

Table 2:

White Paper/ Sites	User Account Model	Authentication Method	Authorization Method	Advantage/Disadvantage
Account Allocations on the Grid, University of Michigan	Pool of template accounts, which creates temporary persistent bindings b/w template account and Grid users	X.509 Certificate	Using local authorization mechanism	Advantage: <ul style="list-style-type: none"> • Allows to map users to individual account templates within a binding pool domain to set different system privileges • Easy maintenance of user accounts Disadvantage: <ul style="list-style-type: none"> • Needs new mechanism to map the resource usage information back to the original user.

Table 2:

White Paper/ Sites	User Account Model	Authentication Method	Authorization Method	Advantage/Disadvantage
Simplifying Administration and Management Processes in the Polish National Cluster	Virtual User Account Systems which consist of following elements: <ul style="list-style-type: none"> • any queuing system • User interface commands to handle queuing system • pool of Virtual User Accounts on each machine • Virtual User Manager Account with a program that can find and assign on the of Virtual User Accounts • Virtual User Server responsible for global translation of user identifiers on all machines as well as user accounting and authorizing a user's access to the machine • modules allowing jobs to other queuing systems to be sent 	N/A	User Account Server Daemon which keeps information about mapping from real users to Virtual User Account identifiers on all systems	Advantage: <ul style="list-style-type: none"> • Simplify user account administration Disadvantage: <ul style="list-style-type: none"> • Need to be maintain list of users who are allowed to use distributed queues which can send tasks to remote machine.

Table 2:

White Paper/ Sites	User Account Model	Authentication Method	Authorization Method	Advantage/Disadvantage
University of Manchester of England	allocate a username from the pool of unleased template accounts For each username in the pool, whether leased or not, an empty file exists in this direc- tory with the same name as the user. When a username is leased, a hard link is made to it, with the same name as the subject. Subject names are stored using the "URL encoding" method - letters and numbers are unchanged, other char- acters are replaced by %HH where HH is their value in hexadecimal. (This is necessary since subjects include slashes and other unfriendly characters.)	x509 certifi- cate	Entry in the gridmap file with the corre- sponding local template or real accounts	Advantage: <ul style="list-style-type: none"> Allows to map users to indi- vidual account templates within a binding pool domain to set different sys- tem privileges Easy maintenance of user accounts Easy maintenance of grid- map file Disadvantage: <ul style="list-style-type: none"> Needs to new mechanism to map the resource usage information back to the original user.

Current Practices

There is a survey titled "Resource Accounting - Current Practices" conducted by Laura F. McGinnis, Grid Forum Accounts working group member. This survey was conducted to understand the current accounting processes used in different sites and to understand the issues the Grid Distributed Computing will introduce. This is the summarization she concluded from the survey.

This survey was conducted so that the Distributed Accounting Working Group(DAWG) would have a more solid context within which to work. Although the sites that responded are generally small, they are a diverse cross-section of the high-performance computing world. The DAWG will be well-served as it works on recommendations and standards to keep in mind the needs of these sites.

Conclusions

- Points of commonality
 - All sites support a variety of resources already.
 - All resource requests are reviewed before they are granted.

- Users must have an account on the specific resource to be permitted to access it.
- Usage must be reported to the site's sponsoring or funding organizations.
- Points of differentiation
 - Not all requests are reviewed on-site. As a result, sites do not have full autonomy to decide who may access their resources.
 - Review criteria are mission-specific and therefore vary from site to site.
 - The review cycle varies from days to months, even at a given site. Some requests might not be reviewed for up to 6 months; others may be granted access in less than a week.
 - Not all sites accept requests from commercial organizations.
 - Usage is calculated differently at all sites.
 - All sites handle overdrafts differently.
 - Sites are accountable to their sponsors, but on different reporting cycles.
 - Different site sponsors require different utilization measures from each sites.
 - Sites have different methods for keeping users up to date on their allocation status.

Not all points of commonality are supportive of GRID computing, but neither are all points of differentiation necessarily barriers. One of the biggest challenges will be addressing the common requirement that users have local accounts to use resources. Local accounts are the basis for usage accounting, which is the primary feedback mechanism between the sites and their funding agencies. On the other hand, variations on how sites review request or calculate usage may allow the sites to maintain their unique identities in the GRID environment, enabling users to make more informed decisions when requesting GRID resources.

More detail information on survey is available on the paper, "Resource Accounting-Current Practices" posted on URL(<http://www.nas.nasa.gov/thigpen/accounts-wg>).

Grid Computing survey results

We have conducted a a survey on Distributed Accounting of the Grid among the Grid Forum Accounts working group. This survey was conducted to collect input from other Grid sites on the types of mechanisms that are deployed and how they are accomplishing distributed accounting. Also to have a better understanding of issues, concerns, or "lessons-learned" that the account-less accounting may have in a Grid environment.

We have received just one response from the Monash University. Their focus is on Economic Models and Resource Brokering system called Nimrod-G(<http://www.csse.monash.edu.au/~raj-kumar/ecogrid/>) Their objective is to implement the resource management and scheduling system that interface with the Grid Middlewares. Currently they require an individual account on each system and uses X509 certificate proxy to authenticate the user. However, they will support any accounts model and authentication method middleware will support. Their resource usage information is provided using Nimrod-G, which provides application level accounting. They have suggested using a "Grid Bank and model" which can interface with the QBank, A CPU Allocations Bank developed by Scott Jackson at the Pacific Northwest National Laboratory.(<http://www.emsl.pnl.gov:2080/docs/mscf/qbank-2.8>)

We are not able to do the complete summarization on the survey questions we have sent out due to lack of response and interest. However, here are list of things that needs to be addressed before the successful account-less model can be implemented. This list will grow as we receive the response from more people.

- Each site needs to determine the mechanism to track the usage information because the Grid middleware doesn't track the resource usage information.
- Needs the mechanism to identify jobs submitted through Grid middleware from the local job-manager's log file.
- Needs the mechanism to post the resource usage information in common defined format for users to access
- Needs to identify common mechanism to authenticate the users.

Distributed Grid Accounting Model

In order for the Grid to scale in the large community, there need to be a new mechanism to authorize users to use the local resources without requiring users to have an account on the local system. Our goal is to design a Distributed Accounting model that can be easily be adapted to the local Account Management System and to the dynamic changes of the Grid environment. Here is the proposed model for the Grid Accounting system at NAS.

Model 1

This model is based on collecting user's distinguish name(DN), jobid, template account name in a permanent Grid accounting log file. Advantage of using this model is that there is no dependency on GIS server to post the job information, so the resource information can be tracked from the local jobmanagers log file. However, this model requires the modification in the Globus GRAM source code to save the information in to a accounting log file. Also, this model doesn't have the capability to track the usage information for the fork jobs.

- Set up template accounts on each of the Grid Resource systems (You can create the template accounts with different system privilege levels, if needed)
- The Grid users are authenticated using the X509 certificate proxy.
- After the authentication is done, grid user will be bind to a template accounts.
- The Grid user will be authorized to use the resources using the template accounts privilege.
- After the successful job submission, Grid users DN, jobid, and template account will be saved in to a grid accounting file.
- After the completion of the job, the resource usage information will be retrieved from the local jobmanger's log file using the jobid and template account save in the grid accounting file.
- The resource usage information will be mailed to the Grid user and will be posted on the GIS server for later references for users.
- The resource usage information will be tracked in our local account management system.

Currently IPG accounting is done using the job log information collected from the 'jobd' program which runs on IPG systems. This method has sole dependency on GIS server been able to publish jobs correctly. With the current release of the GIS server, it's not a reliable method and it has been very difficult task to track the problems of jobs not getting publish. See the following for more detail explanation on the current IPG accounting system.(<http://www.nas.nasa.gov/Groups/Database/ipg.html>)